SOUNDSCAPES

AN-002
APPLICATION NOTE

MAKING GOOD MEASUREMENTS
LEARNING TO RECOGNIZE AND AVOID DISTORTION

by Langston Holland - info@audiomatica.us

INTRODUCTION

The purpose of our measurements is to acquire data that is useful for characterizing the device under test (DUT)[1]. Errors in the measurement system will be added to the behavior of the DUT and obscure what we wish to study. If we want to build a house with walls of a given height, yet use a measurement tape with a false scale, we will fail in our endeavor. Thus it is important that our measurements describe DUT behavior without skew from the measurement system or environment (acoustic or otherwise). When the DUT is an acoustic transducer such as a loudspeaker or microphone, the environment often provides a significant portion of the response. Proper measurement technique involves maximizing the DUT's output (signal) while minimizing the contribution from its environment (noise). There are practical limits as to how far we can go in the pursuit of either goal and this application note primarily concerns the limits imposed on the former; maximizing the desired signal.

There are two types of distortions that we are likely to encounter; linear and nonlinear. The first are relative amplitude and phase variations in the frequency response of a system that do not change over time or drive level and do not add frequencies that were not in the original signal. Linear distortions are often intended, such as when produced by an external equalizer or CLIO’s built-in equalization tools. The second type of distortion is far more likely to cause problems in our measurements. Nonlinear distortion occurs when the system produces frequency content that was not part of the original signal and is difficult or impossible to correct in post-processing. In summary; for any measurement system to produce reliable data, the system must remain within its scale (linear) and consistent in behavior over time (time-invariant). The combined concepts of linearity and time-invariance are known as LTI.

The greater the proportion of nonlinearities present in the system, the less we can trust our measurements to provide accurate data. The most common nonlinearity we will need to watch out for is the production of new frequencies called harmonic distortion. Harmonic means these new frequencies have a fixed relationship to the original (fundamental) frequencies that were included in the stimulus. These relationships tend to be even and odd multiples of the fundamental frequencies, though other new frequencies can occur within these multiples as well, such as with intermodulation distortion.

[1] While this abbreviation is true, the concept of measurement is more correctly seen as a three step process in which the first and last steps, while implicit, are absolutely necessary. The first involves modeling the DUT parameters of interest. In other words, we must have a good understanding of the device or system we are measuring in order to properly interpret the results of our measurements. This also implies being able to predict what the DUT parameters should look like and will guide us in selecting the proper measurements to execute. The second is actually making the measurement. The final step is the analysis and interpretation that yield the estimation of the DUT parameters of interest. This is where we attempt to answer our questions.

From Max Planck, the brilliant theoretical physicist whose work on quantum theory won him the Nobel Prize in Physics in 1918: *Before an experiment can be performed, it must be planned - the question to nature must be formulated before being posed. Before the result of a measurement can be used, it must be interpreted - nature's answer must be understood properly.*

SYSTEM EXAMPLE

In the diagram below we have a typical sound system in which we are attempting to measure a loudspeaker using the LogChirp module. We want to verify proper function and make adjustments to EQ and delay as needed to improve its performance within the room.

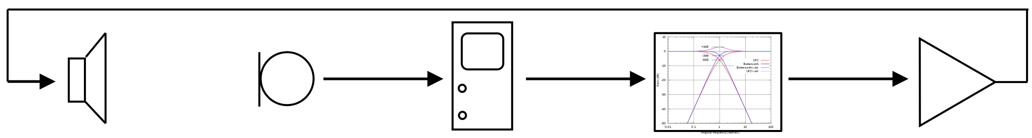

Fig 1: Loudspeaker (DUT), microphone, CLIO, digital processor and amplifier.

In this example we will assume each component in the system is in good working order, thus our first task will be to optimize the measurement system’s gain structure in order to record the highest undistorted signal level with the least possible noise. Our microphone will comfortably handle the SPL it will receive from the loudspeaker and we have adjusted the microphone preamp gain in CLIO so that it shows a good level in the green area of the input meter. The amplifier’s gain has been turned all the way down and CLIO’s output level has been adjusted to show good levels on the input meters of the loudspeaker processor. Now we carefully turn the gain of the amplifier up until we achieve the SPL from the loudspeaker that we know is comfortably within its specifications.

After executing the first measurement we see we have a problem[2] (Fig 2). The frequency response of the loudspeaker seems as expected, but the energy time curve (ETC) of the impulse response (IR) is showing something that seems impossible - energy arrivals earlier than (to the left of) the main arrival from the loudspeaker!

Living in a world governed by causality, we know that we cannot have energy arrivals from the loudspeaker before it is stimulated, so we start a process of elimination and disconnect the cables from the CLIO audio interface to perform a simple loopback measurement and find everything in order. Next we add the loudspeaker processor to the measurement loop and remeasure at the same stimulus output level, +7dBu in this case.

Now we find what seems to be our problem (Fig 3) - the same additional arrivals at the same levels and spacings relative to the main arrival, but now that the simulated time-of-flight has been removed the additional arrivals have moved to the back of the IR record. It’s almost as if the IR was a circular tape and as the main arrival was pushed to the left toward the front of the record, what had previously appeared as pre-arrivals scrolled around to the rear!

[2] For those experienced with room ETC's, you'll notice that the noise floor in Fig 2 is far too low to be a real room. This is a measurement of the processor alone to illustrate the text with greater clarity. An example of a real room impulse response with two "pre-arrivals": http://audiomatica.us/resources/Education/IRdist.jpg. In this situation the voltage output from the measurement microphone exceeded the input capability of the audio interface. Reducing the microphone preamp gain removed these "pre-arrivals" in subsequent measurements.

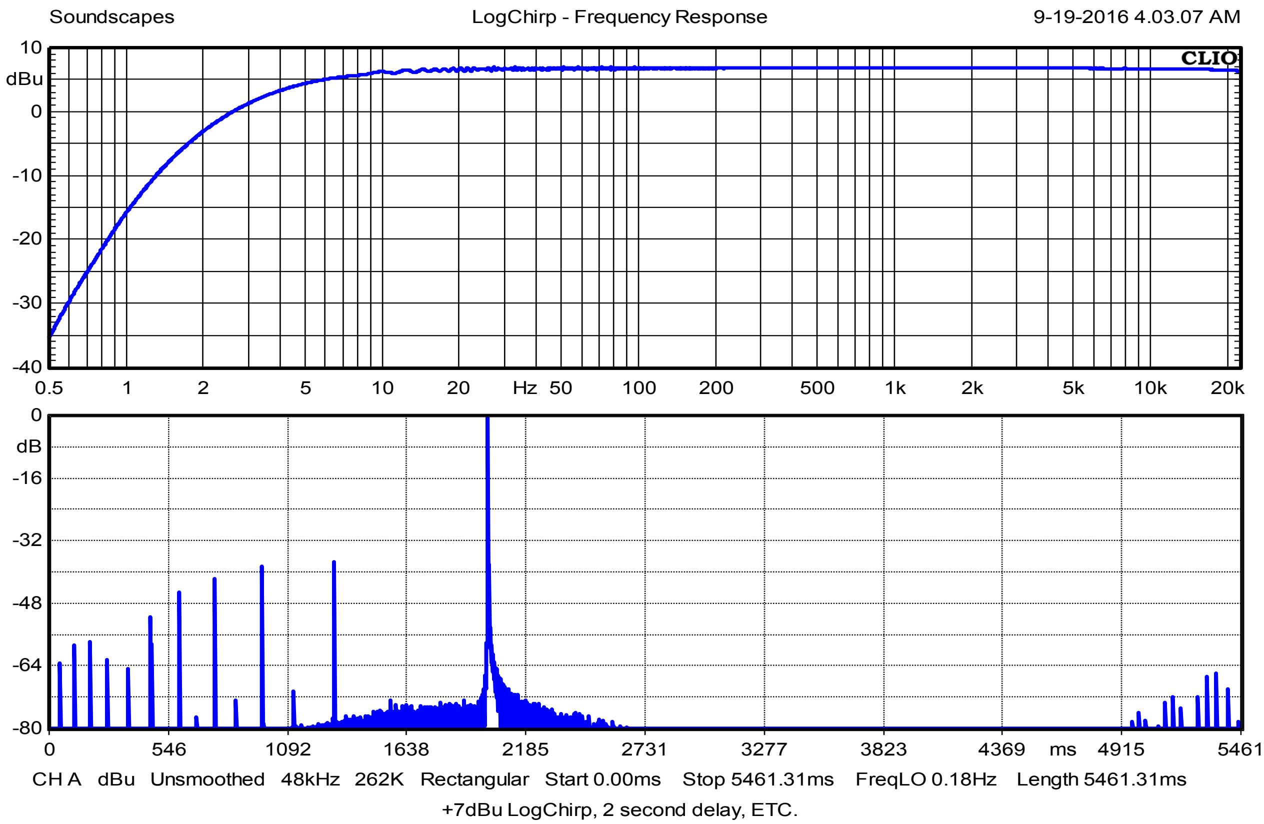

Fig 2: Energy Time Curve showing multiple "arrivals" prior to main signal arrival delayed 2 seconds.

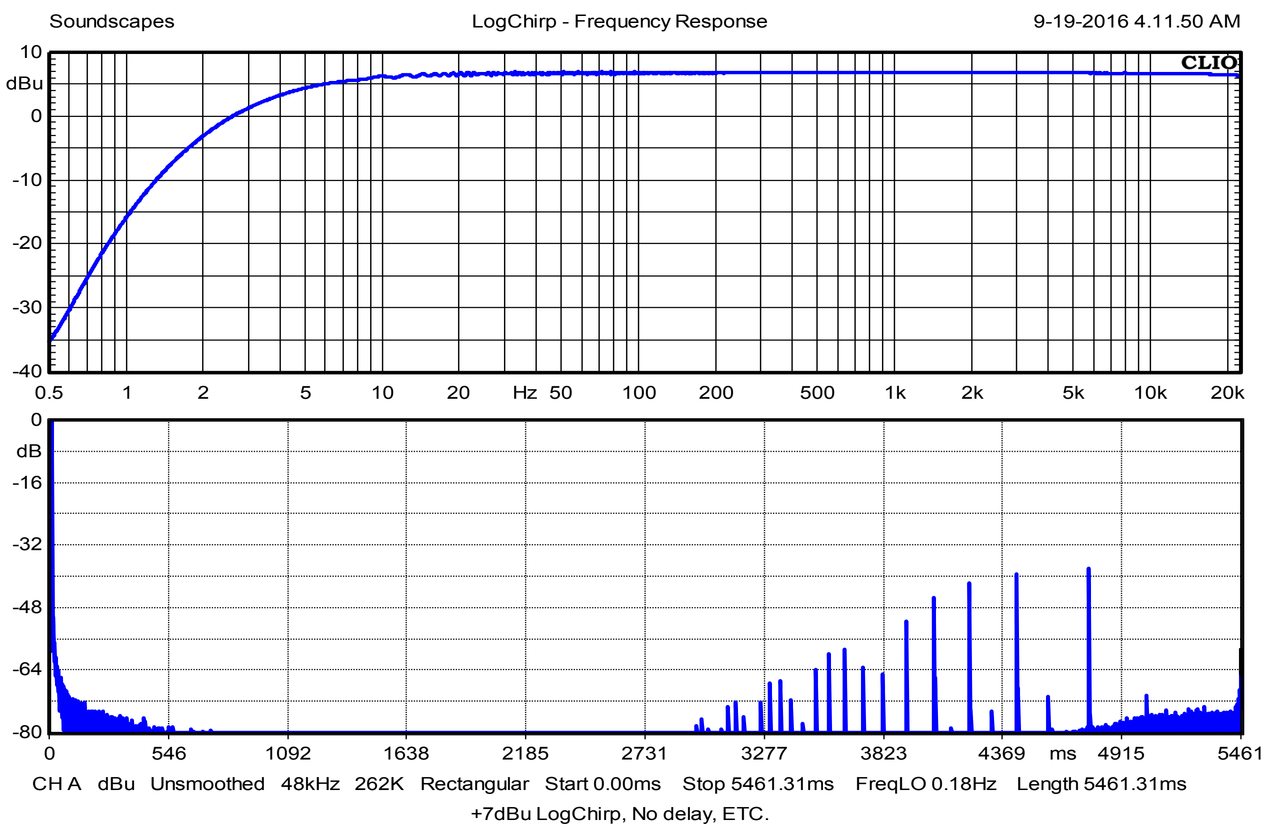

Fig 3: Energy Time Curve showing multiple "arrivals" behind the main signal arrival with delay removed.

Now we suspect our loudspeaker processor may be malfunctioning, but just to be sure we remeasure at a reduced level of 0dBu and we see a perfectly normal response without the additional energy peaks (Fig 4).

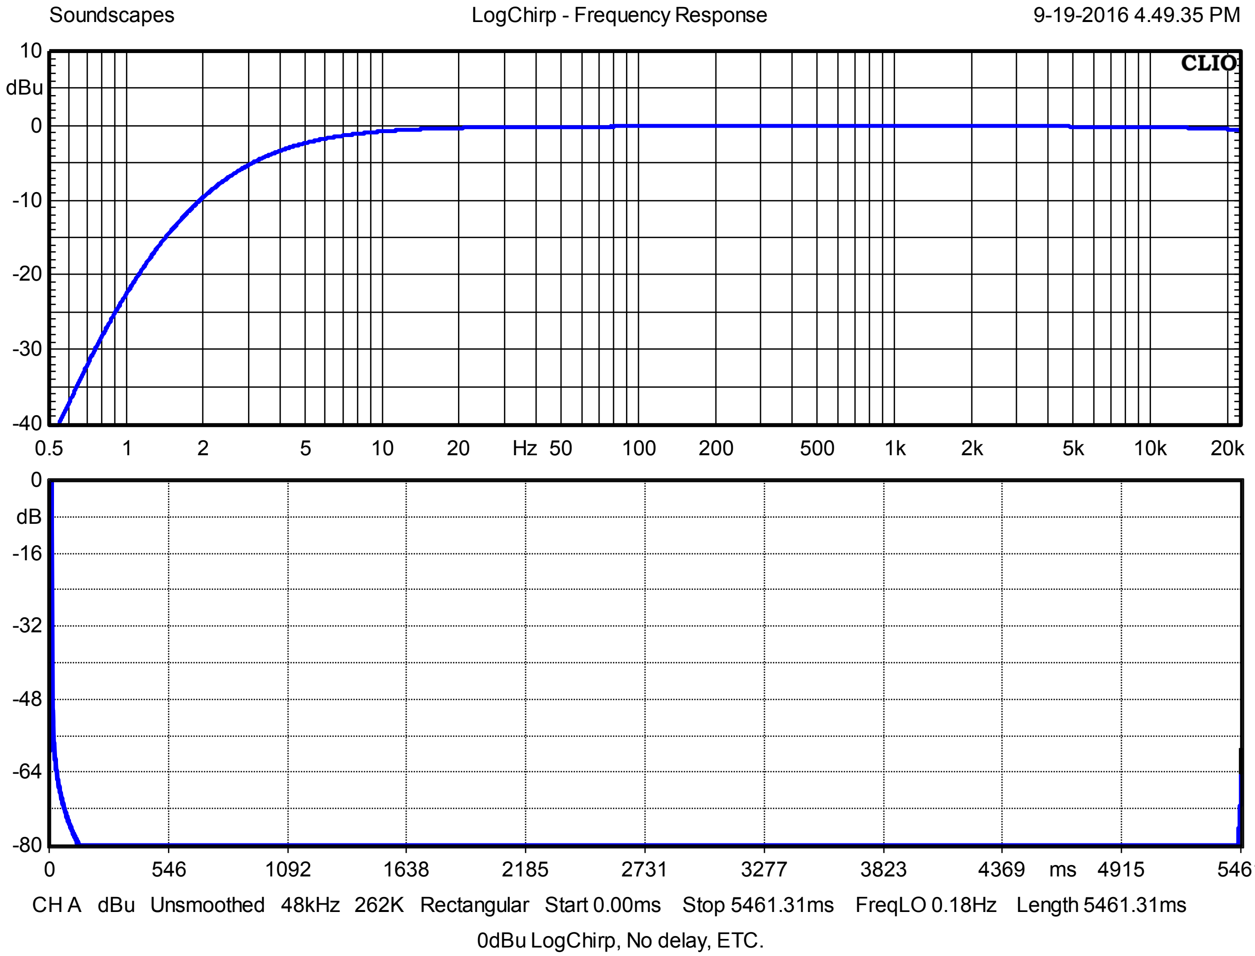

Fig 4: Energy Time Curve at reduced drive level of 0dBu with delay removed.

FFT AND SINUSOIDAL DISTORTION ANALYSIS

Suspecting distortion, we use the FFT module with a 1kHz sine stimulus at the original +7dBu level and see the classic result of harmonic distortion from sine wave clipping (Fig 5). The much greater proportion of odd harmonics (3kHz, 5kHz, 7kHz, etc.) relative to even harmonics (2kHz, 4kHz, 6kHz, etc.) is generally indicative of solid state electronics.

We find in the loudspeaker processor's manual that the meters only monitor its analog stages, so we look into the programming and find gain errors that lead to clipping in the digital domain which did not register on the unit's external metering.

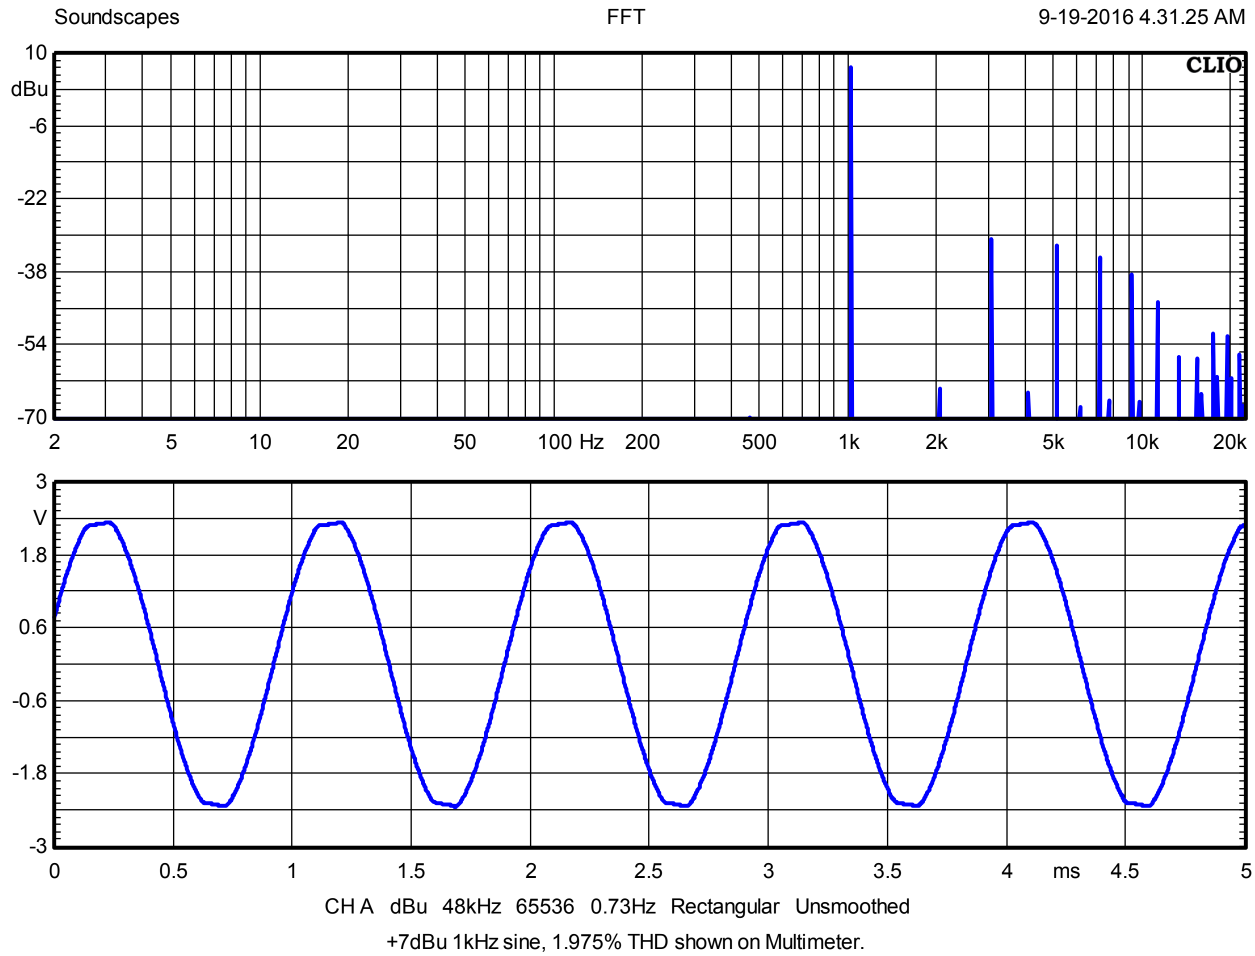

Fig 5: FFT of 1kHz sine at +7dBu drive level.

Further investigation using the Sinusoidal module's powerful distortion analysis (Fig 6) agrees nicely with our FFT module measurement. Distortion traces have been raised 30dB for clarity, thus the THD trace is actually 34dB under the fundamental magnitude trace, or at about 2% THD $\left[10^{-34/20}\right]$.

One of the unusual advantages of CLIO is that it gives you the ability to measure device and system responses with three very different tools. In this application note we use each to give us a better understanding of the distortion that can affect our measurements. The more traditional LogChirp/MLS and FFT modules balance sophisticated post-processing analysis in the former with a real-time, stimulus independent single or dual channel approach in the latter. The Sinusoidal module offers one of the most robust nonlinear signal analysis systems available with its synchronized tracking filter. Nonlinear behavior such as harmonic distortion can be analyzed with great precision using continuous, stepped, or stepped/gated stimuli. The stepped/gated stimulus is one of the few techniques available that is able to accurately measure a device's full bandwidth peak output without the threat of thermal damage.

Comparing the distortion components of the FFT plot in Fig 5 with those from the Sinusoidal module in Fig 6 show excellent agreement. For example, the level difference between the fundamental at 1kHz and the 3rd distortion harmonic is 38dB on both plots. Now for something really interesting, look back at the LogChirp plot in Fig 2. Notice that the closest high level energy peak to the main arrival is also down 38dB. Keep reading to find out why!

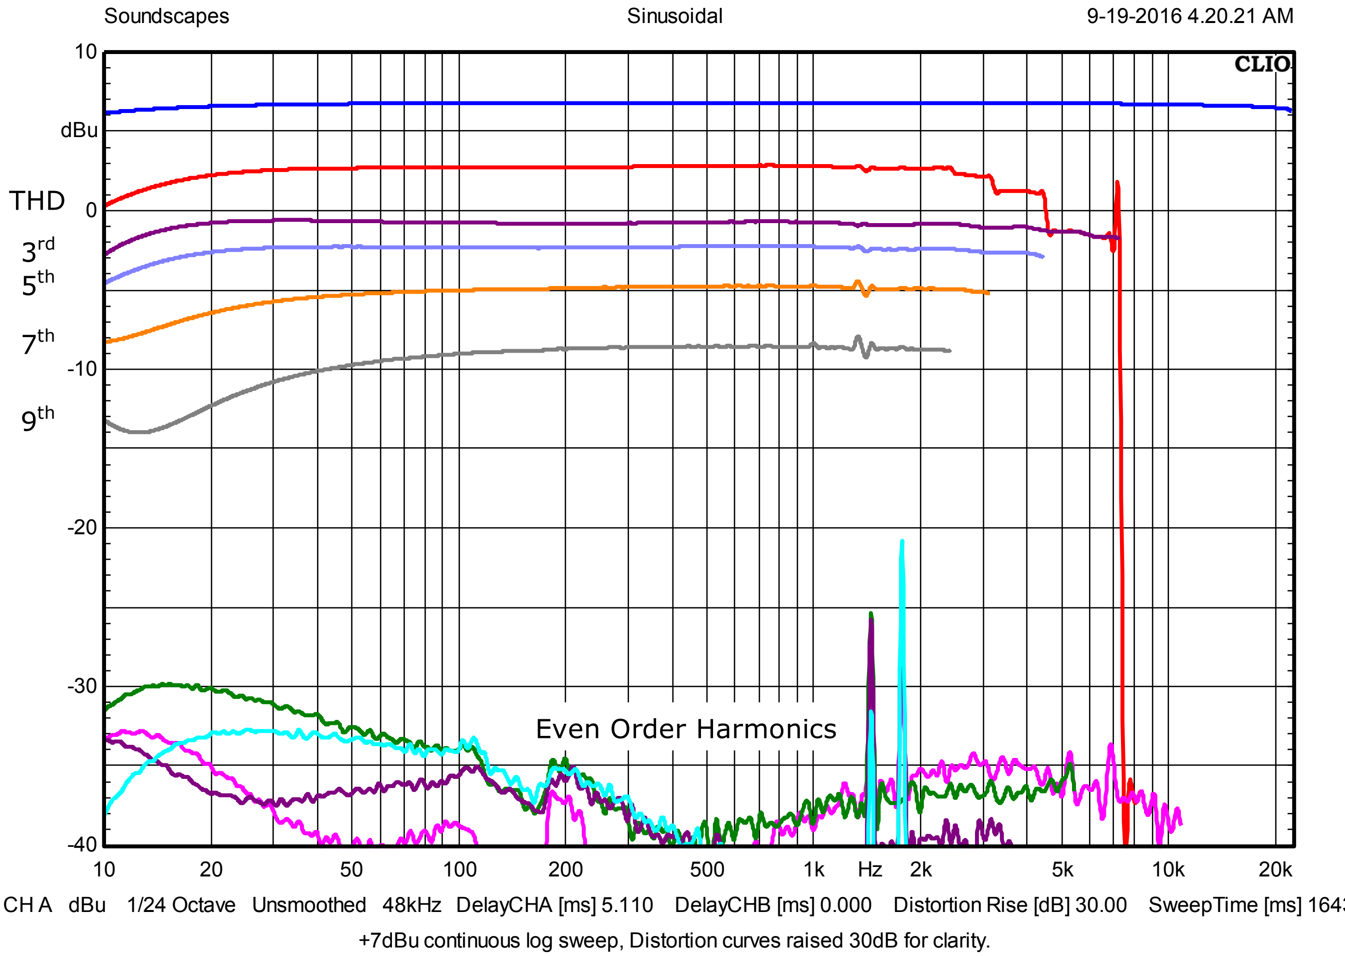

Fig 6: Sinusoidal module slow log sweep at +7dBu drive level with THD and harmonics raised 30dB for clarity.

LOGCHIRP DISTORTION CONCEPTS

Using the ETC's very sensitive vertical log scale, we have seen that if the initial energy arrival is delayed long enough, any distortion products will begin to appear to the left of the main arrival. This is due to the FFT procedure called circular deconvolution used with the periodic log swept sine stimulus. The initial arrival and its distortion products will be separated by an exact number of samples as defined by the sample rate and stimulus start and stop frequencies. The cyclical FFT process will force the distortion components to wrap around to the back of the IR record if there are an insufficient number of samples available between time zero and the initial energy arrival.

Keep in mind that the log swept sine is the only stimulus available to us that distributes the individual distortion harmonics into distinct sample positions. This is one of its great advantages[3] and typically allows us to remove most of the distortion with a post-process window. Other stimuli such as MLS, pink noise and linear sine sweeps smear the distortion components throughout time in the IR record as an additional form of noise that compromises the entire measurement.

[3] For a full treatment on the advantages of the log swept sine stimulus, buckle your seat belt and click here: http://audiomatica.us/resources/Education/IR_Stimuli_Compare.pdf

NEGATIVE TIME

When discussing distortion harmonics separating out of a swept sine measurement in "negative time", we apply this terminology whether the components are to the left of the initial energy arrival or to the right near the end of the IR record. Where they are displayed is merely a function of the number of samples available before the initial energy arrival and the number of samples the FFT process shifts each frequency.

The term negative time comes from the fact that in a typical LF to HF sine sweep, the HF distortion harmonics appear to the measurement math before those frequencies were scheduled to appear. Thus the math dutifully records these harmonics as showing up early (in negative time) relative to the swept fundamentals.

EXAMPLE

You're using the LogChirp module with the log swept sine stimulus that sweeps from LF to HF and you accidentally drive an audio device into clipping. The log sweep stimulus sweeps much more slowly at LF than HF[4]. Say it takes 10ms from the start of the sweep to get to 100Hz and another 15ms for a total of 25ms to get to 200Hz. The goal of the math is to move all of the frequencies to a single impulse at "time zero". Thus, when the math sees 100Hz, it moves it -10ms, when it sees 200Hz, it moves it -25ms, etc. This works perfectly when only the fundamental frequencies show up as expected.

Now say the 100Hz fundamental clips and arrives with the 2nd distortion harmonic of 200Hz, the math sees that 200Hz and moves it -25ms. The problem is that this instance of 200Hz showed up at the same time as the 100Hz fundamental, thus the 2nd harmonic distortion component is placed 15ms prior to the main arrival in "negative time". The 3rd harmonic of the 100Hz fundamental, if it exists, will be placed even further to the left of the main arrival in "negative time", etc.

CONCLUSION

To record measurements with the highest possible S/N ratio, drive the system as hard as possible without clipping anything, including CLIO's audio interface. Keep the measurement setup as simple as possible by bypassing all unnecessary devices and cables in the system. Make sure no compressors or limiters are engaged during the measurement as their behavior is nonlinear by design. When using the LogChirp module with the log sweep stimulus, make a habit of inspecting the ETC to verify the absence of clipping.

ACKNOWLEDGEMENT

I want to thank Daniele Ponteggia of Audiomatica s.r.l. for his many corrections and insights during the writing of this paper. Remaining errors, if any, are the fault of the author. Comments solicited.

REFERENCES

[1] S. Temme, "Audio Distortion Measurements", presented at the AES 11th International Audio Test and Measurement Conference, *J. Audio Eng. Soc.,* (Portland, OR, 1992 May).

[2] A. Farina, "Simultaneous measurements of impulse response and distortion with a swept-sine technique", presented at the 108th Convention of the Audio Engineering Society, *J. Audio Eng. Soc. (Abstracts),* (Paris, France, 2000 Feb), preprint 5093.

[3] C. T. Tan, B. C. J. Moore and N. Zacharov, "The Effect of Nonlinear Distortion on the Perceived Quality of Music and Speech Signals", *J. Audio Eng. Soc.,* Vol. 51, No. 11, pp. 1012-1031 (2003 Nov).

---

[4] Otherwise known as a pink spectrum. The log sweep provides a significant advantage over linear sweeps in maximizing the signal to noise ratio because acoustic environments generally have far more noise at LF than HF. Having the stimulus spend more time at lower frequencies adds more signal to the measurement process where it is needed.